**<u>Amendments to the Drawings:</u>**

The attached replacement sheet of drawings is a replacement sheet for FIG. 4. In view of the Examiner's remarks, it was noticed that the formal drawing sheet including FIG. 4 did not show the different relative sizes of the boxes labeled "CACHE RAM CARDS" and "NETWORK INTERFACE CARDS" as shown in the informal drawings originally filed with the application. Therefore, in the new attached FIG. 4 replacement sheet shows the boxes of different relative sizes corresponding to the sizes shown in the informal drawings as originally filed.

Attachment: Replacement Sheet 3/16 (FIG. 4)

Annotated Sheet Showing Changes 3/16 (FIG. 4)

**REMARKS/ARGUMENTS**

By this amendment, claims 27-30 have been cancelled. Claims 2-9, 11-14, 16-23, and 25-26 remain in the application.

By this amendment, a replacement sheet for FIG. 4 is being submitted to conform the formal drawing for FIG. 4 to the informal drawing as originally filed. As described on page 13 line 15 to page 14 line 10 of the applicants' specification as originally filed, the data movers servicing the higher ranking (i.e., the most popular ) movies have fewer cache RAM cards 41 and more network interface cards 42 than the data movers serving the lower ranking (i.e., less popular) movies.

In paragraph 3 on pages 6-7 of the final Official Action, claims 27-30 were objected to under 37 CFR 1.75(c) as being of improper dependent form for failing to further limit the subject matter of a previous claim. In response, claims 27-30 have been cancelled.

In paragraph 5 on page 7 of the final Official Action, claims 2-9, 11-14, 16-23, and 25-30 were rejected were rejected under 35 U.S.C. 103(a) as being unpatentable over Armstrong et al. (WO 2000/60861) in view of Mizutani (U.S. Patent 6,115,740). In response, applicants respectfully traverse on the grounds that the claimed invention would not have been obvious from Armstrong and Mizutani.

The invention is directed to a video file server and method of operating the video file server for providing video-on-demand access to movies. The video file server includes a cached

disk storage system including a primary cache and disk storage for storing the movies, and a multiplicity of data mover computers coupled to the cached disk storage system for streaming video data from the cached disk storage system to clients in the data network. Each of the data mover computers has a local cache. This is shown in applicants' FIG. 1 as reproduced below and described in applicants' specification on page 7 line 12 et seq.:

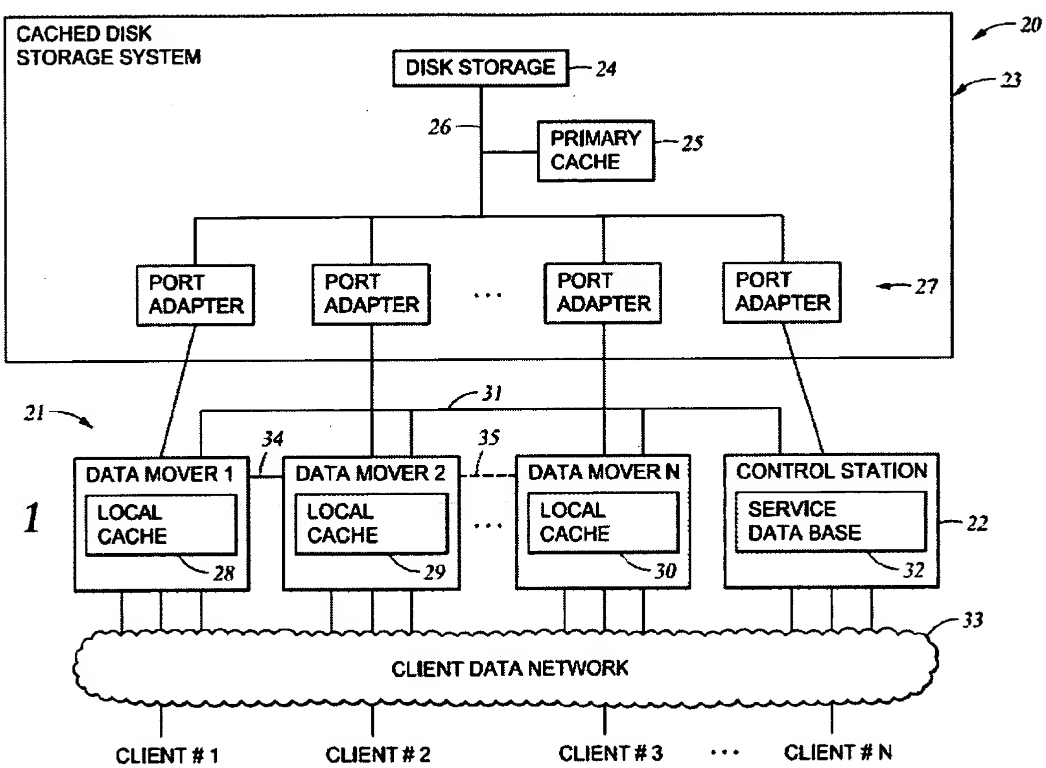


Fig. 1

The applicants' independent claims 2 and 16 are more specifically directed to ranking the movies with respect to popularity, and pre-assigning a respective set of the data movers for

servicing video streams for each movie ranking, wherein the data movers in the respective sets of data movers are configured differently for providing more network interface resources for very popular movies and for providing more local cache memory resources for less popular movies.

The ranking of movies with respect to popularity is shown in applicants' FIG. 5, reproduced below:

| MOVIE RANKING | MOVIE TITLE | FREQUENCY OF ACCESS | DATA MOVER SET | NO. OF RESERVED VIDEO STREAMS | NO. OF ACTIVE VIDEO STREAMS | LOCATION ON DISK | ANY LOCATION IN CACHE |
|---|---|---|---|---|---|---|---|
| 1 | — | 25,000 | 1 | 5,000 | 3,497 | — | — |
| 2 | — | 11,000 | 2 | 2,200 | 1,738 | — | — |
| 3 | — | 7,900 | 3 | 1,580 | 2,300 | — | — |
| 4 | — | 6,900 | 3 | 1,380 | 790 | — | — |
| 5 | — | 4,000 | 4 | 800 | 570 | — | — |
| 6 | — | 2,600 | 4 | 520 | 640 | — | — |
| 7 | — | 2,300 | 4 | 460 | 198 | — | — |
| 8 | — | 2,000 | 4 | 400 | 340 | — | — |
| 9 | — | 1,700 | 5 | 340 | 270 | — | — |

*Fig. 5*

The process of ranking the movies with respect to popularity is described on page 11 lines 10-19 of applicants' original specification as follows:

> When a movie is first released for VOD distribution, it will have an industry rating, such as the "Blockbuster" rating, with respect to other popular movies. Therefore, the newly released movie can be ranked with respect to the popular movies already stored in the video file server, and the rank of each lower

rated movie in the video file server can be decreased by one level as the newly released movie is written to disk storage of the video file server. New movie releases, for example, are written into the disk storage at the time of day of minimum demand, for example, about 3:00 a.m. Each day, the number of accesses for each movie is recorded to compute a running average of the frequency of access of each movie, in order to re-adjust the rank of each movie, and make effective the adjusted rank at the time of day of minimum demand.

The pre-assigning of a respective set of data movers for servicing video streams for each movie ranking is shown in applicants' FIG. 3, reproduced below:

| MOVIE RANKING | DATA MOVER SET |
|---|---|
| 1 | SET 1 = {DM1, DM2, DM3} |
| 2 | SET 2 = {DM3, DM4} |
| 3 | SET 3 = {DM5} |
| 4 | SET 3 = {DM6} |
| 5 | SET 4 = {DM7} |
| 6 | SET 4 = {DM7} |
| 7 | SET 5 = {DM8} |
| 8 | SET 5 = {DM8} |
| ⋮ | ⋮ |

*Fig. 3*

The pre-assigning of a respective set of data movers for servicing video streams for each movie ranking is described on page 11 line 20 to page 23 line 6 of applicants' original specification, as follows:

> In accordance with one aspect of the present invention, for each movie ranking, a particular set of data mover resources are preassigned for servicing client access to the movie having the movie ranking. With reference to FIG. 3, for example, there is shown a table associating each movie ranking with a data mover set pre-assigned to service the movie having the associated ranking. Therefore, there is a certain maximum number of video streams that are available to each movie, corresponding to the total number of video streams that all of the data movers in the set can supply to the client data network. There is also a certain maximum amount of data mover local cache memory capacity available for servicing each movie, corresponding to the total amount of local cache memory in all of the data movers in the set.

The respective sets of data movers being configured differently for providing more network interface resources for very popular movies and for providing more local cache memory resources for less popular movies was shown in applicants' FIG. 4 as originally filed and as described on page 13 lines 15-23 of applicants' original specification, as follows:

> As shown in FIG. 4, the physical configuration of each of the data movers 21 depends on the rank of the movie or movies to be serviced by the data mover. In particular, the data movers servicing the higher ranking (i.e., the most popular) movies have fewer cache RAM cards 41 and more network interface cards 42 than the data movers servicing the lower ranking (i.e., less popular) movies. This is a consequence of the fact that the cache RAM has sufficiently high bandwidth

that a large number of video streams can be serviced from one copy of a movie in the cache RAM. For less popular movies, not all of this bandwidth can be utilized, so that the ratio of network interface cards to cache RAM cards will fall for less popular movies.

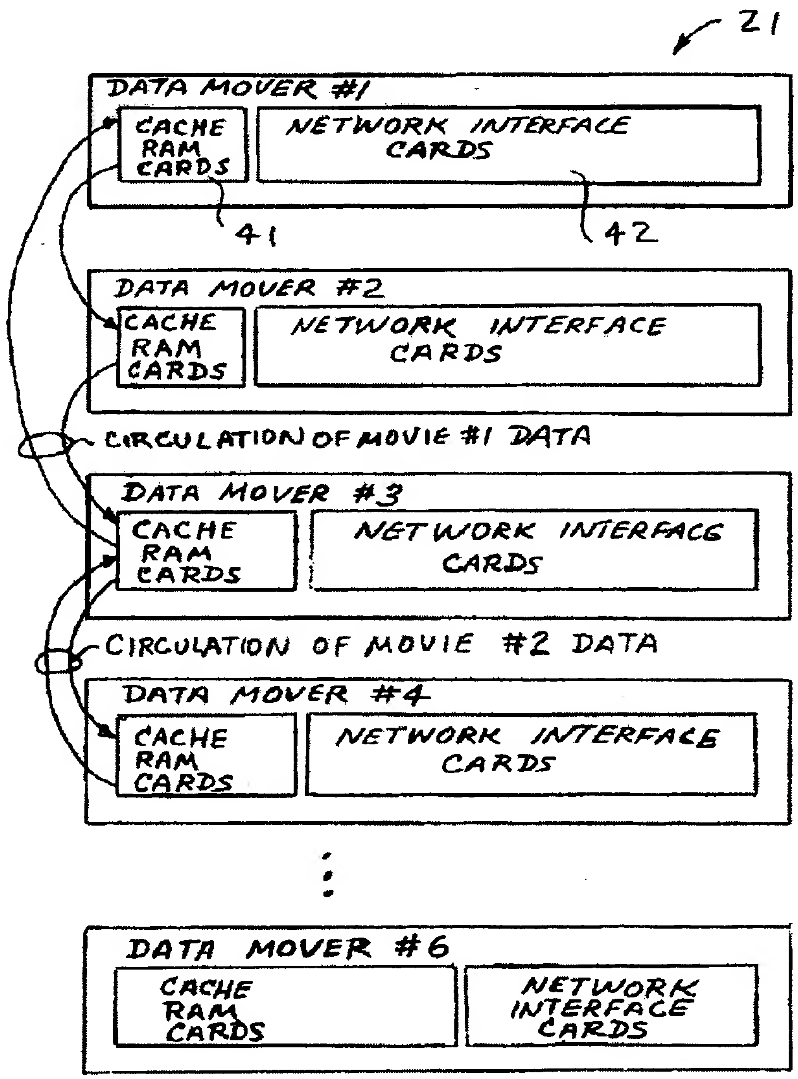


FIG. 4

The policy of the Patent and Trademark Office has been to follow in each and every case the standard of patentability enunciated by the Supreme Court in Graham v. John Deere Co., 148 U.S.P.Q. 459 (1966). M.P.E.P. § 2141. As stated by the Supreme Court:

> Under § 103, the scope and content of the prior art are to be determined; differences between the prior art and the claims at issue are to be ascertained; and the level of ordinary skill in the pertinent art resolved. Against this background, the obviousness or nonobviousness of the subject matter is determined. Such secondary considerations as commercial success, long felt but unsolved needs, failure of others, etc., might be utilized to give light to the circumstances surrounding the origin of the subject matter sought to be patented. As indicia of obviousness or nonobviousness, these inquiries may have relevancy.

148 U.S.P.Q. at 467.

The problem that the inventor is trying to solve must be considered in determining whether or not the invention would have been obvious. The invention as a whole embraces the structure, properties and problems it solves. In re Wright, 848 F.2d 1216, 1219, 6 U.S.P.Q.2d 1959, 1961 (Fed. Cir. 1988).

Armstrong discloses a method and apparatus for hierarchical distribution of video content for an interactive information distribution system. (Title.) A first embodiment shown in FIG. 1 comprises a remote server head-end 110R and a plurality of head-ends $110_i$ through $110_n$ each being coupled to at least one of a corresponding plurality of neighborhoods $130_i$ through $130_n$. (Page 6, lines 13-19.) Each head-end comprises a host workstation 112, a video stream server 114, and a primary storage partition 118 comprised of an array of hard drives. (Page 6, lines 20-24.) Video streams are transmitted from the video stream server 114 to the subscriber's respective subscriber

equipment comprised of a set-top box 142, a display 144 and a control device 145. (Page 6 line 32 to page 7 line 2.) The remote server head-end 110R also includes a secondary storage partition 119R coupled to the remote server 114R. (Page 7, lines 11-13.) A primary storage partition 118 of a head-end 110, including the remote server head-end 110R, is used to store frequently requested video assets. Alternately, the secondary storage partition 119 of the remote server head-end 110R is used to store infrequently requested video assets. (Page 7, lines 17-20.) The content manager 120 tracks the number of requests for a video asset and produces an asset request rate. An operator using the host workstation 112 defines a threshold rate for each video asset. The content manager 120 periodically compares the asset request rate against the threshold rate for each video asset in the system 100. If the asset request rate traverses the threshold rate for the video asset, then the video asset is stored on the primary storage partitions 118 and 118R at each of the head-ends 110 and 110R. If the asset request rate does not traverse the threshold rate for a video asset, then the video asset is stored on the secondary storage partition 119R at the remote server head-end 110R. In this manner video assets are dynamically distributed throughout the interactive information distribution system 110. (Page 7, lines 21-31.)

Armstrong discloses a second embodiment in FIG. 2. In FIG. 2, the primary storage device 216 of the head-ends 210 is apportioned into at least two storage partitions designated as a primary storage partition 218, and a secondary storage partition 219. (Page 9, lines 14-16.) The primary storage partition 218 on the primary storage device 216 at each head-end 210 is used to store frequently requested video assets and temporarily cached library video assets. Each primary storage partition 218 at each head-end typically has the same frequently requested video assets as any other

head-end 210. The secondary storage partition 219 is used to store portions of the infrequently requested video assets. An entire library of infrequently requested video assets is divided and stored amongst the plurality of head-ends 210 at each of the secondary storage partitions 219 on their respective primary storage devices 216. An infrequently requested video asset is typically stored on the secondary storage partition 219 at a single head-end 210. However, the request rate for that video asset may warrant additional storage at other head-ends. As such the content may be replicated and stored thereafter. In this manner, video assets that do not warrant storage across the entire system of head-ends 210 in the interactive information distribution system 200, may still be dynamically stored at multiple head-ends 210. Such dynamic storage corresponding to those neighborhoods having hither request rates than others is made in accordance with an algorithm that allows maximum access to the video titles with minimum network cost associated with their delivery. (Page 10, lines 1-19.) A threshold rate is a value for each requested video asset, established by the service provider in the interactive information distribution system 100, which defines a level to be considered as frequent or infrequent request by the subscribers. Each video asset may have multiple threshold rates. Multiple threshold rates are set to establish various parameters for the storage locations of video information. Such parameters include discarding the video asset, storing it as a single head-end 110, replicating the video asset and storing it at more than one head-end 110 where the request rate warrants it, or storing it as all the head-ends 110 across the entire interactive information distribution system 100. (Page 14, lines 1 to 28.)

Page 8 of the final Official Action says: "Armstrong teaches, 'wherein the movies are ranked with respect to popularity ...' by disclosing primary storage partition 218 is used to store frequently requested video assets and secondary storage partition 219 is used to store infrequently requested video assets (page 10, lines 1-10)." Applicants respectfully disagree. Ranking is different from simply classifying movies as either frequently requested or not frequently requested. The plain meaning of the verb "rank" is "To arrange in a series in ascending or descending order of importance." (See, for example, the definition on page 825 of Rudolf F. Graf, Modern Dictionary of Electronics, Butterworth-Heinemann, Newton, Ma 1997.) Such ranking of movies is shown in applicants' FIG. 5.

In response to applicants' argument, page 2 of the final Official Action refers to a definition of "rank" from the Webster's New World Dictionary. This definition, however, is for the word "rank" used as a noun. The applicants' independent claims 1 and 16 are using the word "ranked" as a verb, and further specify "a respective set of the data movers are pre-assigned for servicing video streams for each movie ranking" and "the data movers in the respective sets of data movers are configured differently for providing more network interface resources for very popular movies and for providing more local cache memory resources for less popular movies." The processes defined in applicants' claims 2 and 16 are not simply giving precedence or more resources to more popular movies than less popular movies. The process in applicants' claim 26 is even more detailed: "the method includes ranking the movies with respect to popularity, assigning a respective set of the data movers for servicing video streams for each movie ranking,

and configuring the data movers in the respective sets of data movers differently for providing more network interface resources for very popular movies and for providing more local cache memory resources for less popular movies."

Page 8 of the final Official Action says: "Armstrong teaches, 'wherein the data movers in the respective sets of data movers are configured differently for providing more network interface resources for very popular movies and for providing more local cache memory resources for less popular movies' by disclosing headend $210_2$-$210_n$ comprise primary storage partition 218 is used to store frequently requested video assets and secondary storage partition 219 is used to store infrequently requested video assets." Applicants respectfully disagree. It appears that each head-end data mover in Armstrong (FIGS. 1 and 2) has the same configuration with respect to cache resources and network interface resources. The applicants' data movers (FIG. 4) are configured differently by having fewer cache RAM cards and more network interface cards in the data movers assigned to storing and servicing the more popular movies than in the data movers assigned to storing and servicing the less popular movies. (Applicant's original specification, page 13 lines 15-23.) In addition, the portions of Armstrong reproduced above suggest that more of the local cache memory will be used for storing more popular movies than less popular movies.

In the response to the applicants' argument, page 3 of the final Official Action notes that Armstrong teaches when a subscriber requests to view a frequently requested video asset or "popular movie" that is stored on primary storage partition 218, the movie is immediately delivered to the requesting subscriber. On the other hand, it a subscriber requests to view an infrequently requested video asset or "less popular movie" which is not stored on primary storage partition, 218,

then resources for the less popular movie must be retrieved from other headends with the system. This is true if the less popular movie is not already in the secondary storage partition, but the primary storage partition 218 is not a network interface resource; instead, the primary storage partition 218 is a memory resources, as can be seen since the primary storage 216 is subdivided or partitioned into the primary storage partition 218 for the more popular video assets and the secondary storage partition 219 for the less popular video assets. In the context of applicants' specification (page 10, lines 10-11), the primary storage 216 of a head-end serves as a local cache for the head-end since the more popular movies are "kept" in the local cache by keeping them in the primary storage partition. Nor should a video asset or movie that would be retrieved from a network and stored in cache be considered to be a cache memory resource, since applicants' claims call for "local cache memory resources for less popular movies." More importantly, the applicants' the claims 2, 16, and 26 are not specifying that one part of a data mover services popular movies and another part of the data mover services less popular movies. Instead the applicants' claims 2, 16, and 26 specify that data movers that service more popular movies have fewer cache memory resources and more network interface resources than data movers that service less popular movies. (This is how the Examiner construed these claims on pages 6 to 7 of the final Official Action in his objection to claims 27-30 under 37 C.F.R. 1.75(c).)

Where the prior art references fail to teach a claim limitation, there must be "concrete evidence" in the record to support an obviousness rejection. "Basic knowledge" or "common sense" is insufficient. In re Zurko, 258 F.3d 1379, 1385-86, 59 U.S.P.Q.2d 1693, 1697 (Fed. Cir. 2001).

Page 8 of the final Official Action says: "Armstrong fails to disclose a respective set of data movers pre-assigned for servicing video streams for each movie ranking." The final Official Action cites Mizutani for this feature.

Mizutani (US 6,115,740) discloses a video file server system for dynamically allocating contents and delivering data. (See title.) The video server system has a plurality of video servers having respective contents storing units for storing contents and respective contents delivering units for delivering contents. A management server has a stream supply information managing unit for managing stream supply information relative to the delivery of the contents and a contents dynamic allocating unit for controlling the storage of the contents between the video servers to dynamically allocate contents based on stream supply information from the stream supply information managing unit. (Abstract.) Mizutani says that different kinds of content C0, C1, C2, ... (e.g., different digitally moving image data, col. 1, lines 15-16) can be stored in a video server system. Each video server in the system can deliver a maximum number (Nstrm) of streams. (Col. 1, lines 48-49.) The maximum number of streams of content that can be delivered at one time from an entire video server system may be increased by increasing the number of installed video servers. (Col. 1, lines 30-33.) In order to avoid rejection of a request for the delivery of a content stream C1, it is necessary that the content C1 be stored in the video server beforehand in expectation of access to the content C1. (Col. 1, lines 1-4.) The estimated number of video servers which can be installed (Nvsa) is the sum for all *i* of Pi/Nstrm, where Pi

represents the maximum number of times that each of the contents is simultaneously accessed per day, and *i* represents the type of a content. (Col. 2, lines 9-20.)

Mizutani says it has been customary to predict concentrated access to certain contents, estimate the number of video servers to be installed, and statistically allocate appropriate contents in the video servers before the video system is put into service. (Col. 3, lines 1-5.) Mizutani says that this conventional static contents allocation scheme usually results in an excessive estimate of the required number of servers to be installed (col. 3, lines 36-37) because of an incorrect assumption that the maximum numbers Pi of times that the respective contents Ci are simultaneously accessed occur at a common time (col. 3, lines 6-11). Instead, in normal circumstances, different users access different kinds of contents at different times. For example, news programs are popular in the morning, and movies are popular in the evening, and some users prefer to see video programs early in the evening and others late in the evening. (Column 3, lines 30-35.)

Mizutani's solution to the problems of the conventional static contents allocation scheme (col. 3, lines 26-28) is to dynamically allocate the contents (col. 3, lines 48-56). The contents are dynamically allocated by detecting whether at least the number of streams of a content stored in a video server or the predictable number of accesses exceeds a corresponding threshold value or not, and if the number exceeds the threshold value, controlling the storage of the content between the video server or another video server, for thereby dynamically allocating the content. (Column 4, lines 3-11.) Predicted values used by the video file server system for dynamically

allocating contents include a predicted maximum number A(s,t) of times that a video server *s* is simultaneously accessed at a time *t*. (Col. 6, lines 6-24.) If there is a request from a user at time t, then the video server whose A(s,t) is the smallest serves as a delivering video server for delivering a requested content. (Col. 6, lines 25-27.) Another predicted value is a number B(i,t) of lacking resources of the content i predicted at the time t. (Col. 6, lines 28-29.) The predicted maximum number B(i,t) is periodically checked for all contents, and the contents are dynamically allocated by being copied, moved, and deleted so that B(i,t)=0 as much as possible. (Col. 6, lines 51-54.) Contents are allocated according to a video server having smallest predicted number of simultaneous accesses at a given time. (Col. 14, lines 8-18.)

Pages 8 to 9 of the final Official Action says: "Mizutani teaches using the predicted number of times that the content I is simultaneously accessed at the time t is represented by Pi(t) and the equation for B(i,t) is used to determine if content is lacking resources to determine how many streams on each server are necessary to facilitate requests (Col. 6, lines 32-38)." As introduced above, however, B(i,t) is a number of lacking resources of the content i predicted at the time t (Col. 6, lines 28-29), and ranking is different from a number of lacking resources.

Page 9 of the final Official Action says: "Figure 7 [of Mizutani] further discloses pre-assigning content, C0 and C1, to servers SV0 and SV1 or 'data movers'. Accordingly, it would have been obvious to one of ordinary skill in the art at the time the invention was made to modify Armstrong with the teachings of Mizutani in order to preassign data movers to service video streams for the benefit of making more resources available for more popular content."

However, the applicants' invention of claim 2 does not result merely by assigning copies of popular content to be stored in the local cache of particular server head-ends in FIG. 1 or FIG. 2 of Armstrong. Nor does the prior art as a whole suggest that Armstrong and Mizutani should be combined or modified as required to arrive at the applicants' invention.

As introduced above, Armstrong uses a threshold technique for hierarchical distribution of video content in a video-on-demand system. Armstrong's technique appears entirely suitable for its intended purpose.

Mizutani finds fault with the prior art static contents allocation scheme (col. 3, lines 26-28) and teaches instead a method of dynamically allocating contents. It is respectfully submitted that from the viewpoint of FIG. 16 as a point of origin, Mizutani and the applicants of the present invention go off in different directions in an attempt to provide more efficient allocation of video server resources and thus avoid installation of an excessive number of video servers to satisfy client demand. Mizutani does not appear to care which contents are popular and which are not, because the invention of Mizutani should dynamically move content between the servers to suit changing conditions. A reference such as Mizutani should be considered as a whole, and portions arguing against or teaching away from the claimed invention must be considered. Basch & Lomb, Inc. v. Barnes-Hind/Hydrocurve, Inc., 796 F.2d 443, 230 U.S.P.Q. 416 (Fed. Cir. 1986), cert. denied, 484 U.S. 823 (1987).

In short, neither Armstrong nor Mizutani appear to suggest that the content should be ranked and the video servers in a video-on-demand system should be configured differently for

providing more network interface resources for more popular movies and for providing more local cache memory resources for less popular movies.

With respect to claim 4, for the same reasons, since Armstrong appears entirely suitable for its intended purpose, it is not seen why one of ordinary skill would have been motivated to modify Armstrong in view of FIG. 16 of Mizutani.

With respect to claim 5, page 11 of the final Official Action says: "Armstrong fails to explicitly disclose transferring the movie data to a data mover servicing a next higher/lower movie ranking." It is also not clear whether the links in Armstrong between the head end servers are direct or not. Applicants also disagree with the contention that it is notoriously well known in the art to transfer movie data to servers serving a next higher/lower movie ranking, since the examination should be based on actual evidence, and in the context of applicants' specification, ranking is something more than simply classifying movies as either frequently requested or not frequently requested.

With respect to claim 12, see the discussion above with respect to applicants' claim 2.

With respect to claim 13, see the discussion above with respect to applicants' claim 5.

With respect to claim 16, see the discussion above with respect to applicants' claim 2.

With respect to claim 18, see the discussion above with respect to applicants' claim 4.

With respect to claim 19, see the discussion above with respect to applicants' claim 5.

With respect to claim 26, see the discussion above with respect to applicants' claim 2.

In view of the above, reconsideration is respectfully requested, and early allowance is earnestly solicited.

Respectfully submitted,

Richard C. Auchterlonie
Reg. No. 30,607

NOVAK DRUCE & QUIGG, LLP
1000 Louisiana, 53rd Floor
Houston, TX 77002
713-571-3400

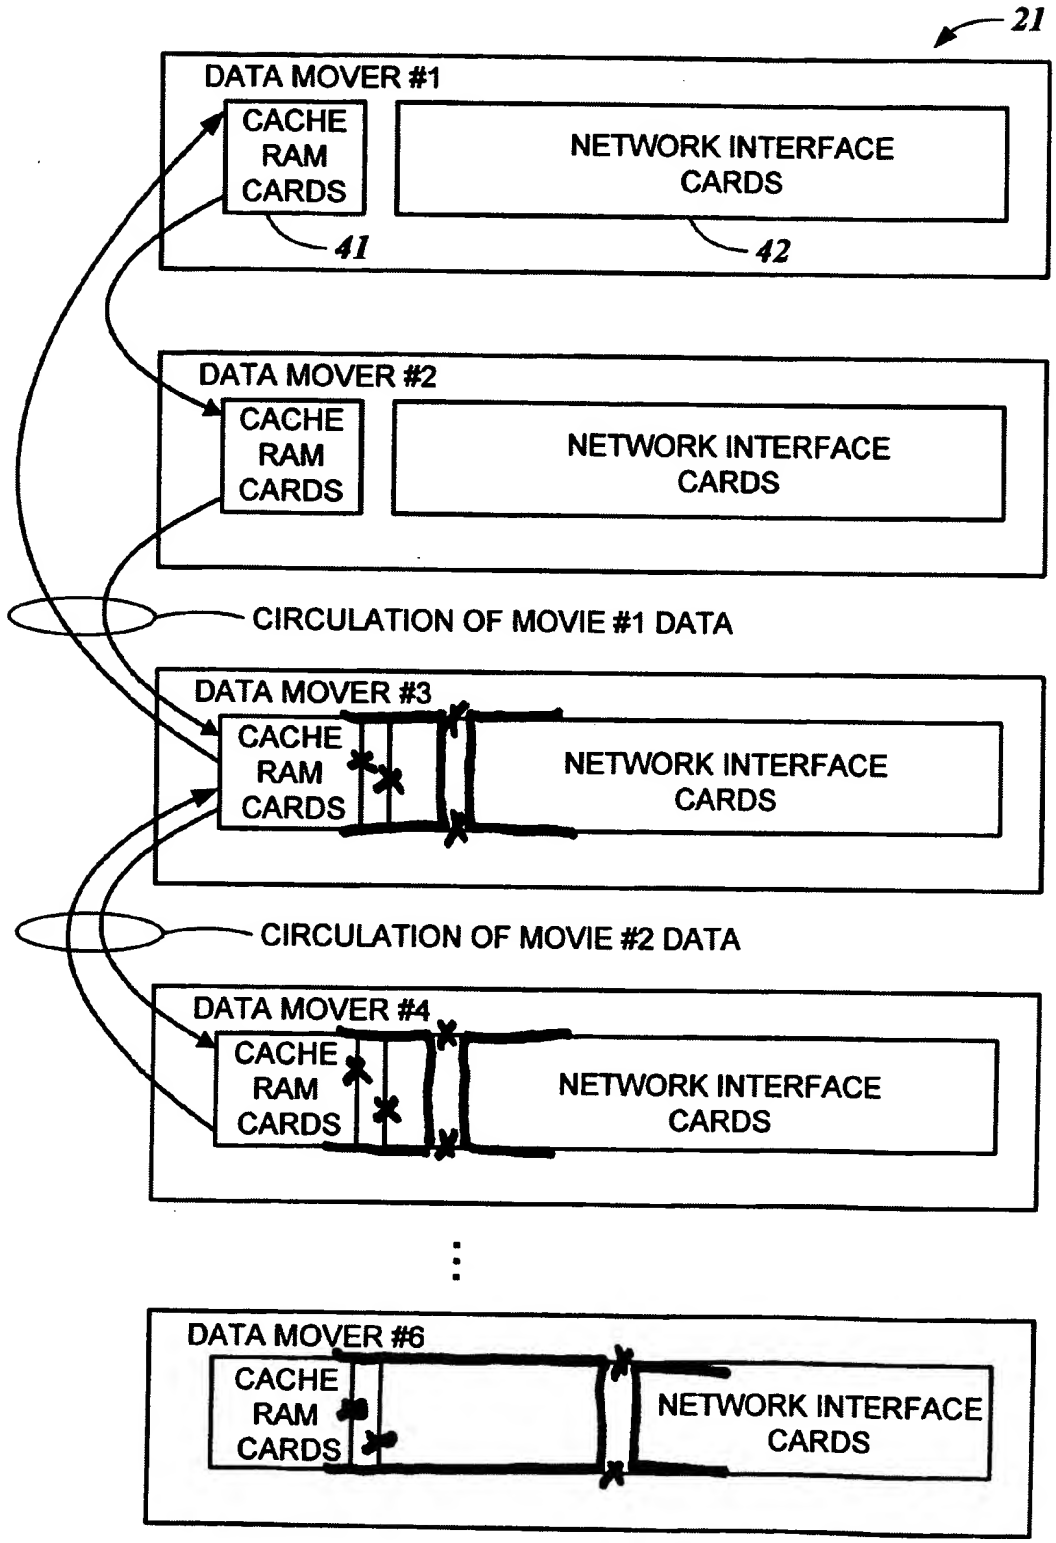
21
DATA MOVER #1
CACHE RAM CARDS
NETWORK INTERFACE CARDS
41
42
DATA MOVER #2
CACHE RAM CARDS
NETWORK INTERFACE CARDS
CIRCULATION OF MOVIE #1 DATA
DATA MOVER #3
CACHE RAM CARDS
NETWORK INTERFACE CARDS
CIRCULATION OF MOVIE #2 DATA
DATA MOVER #4
CACHE RAM CARDS
NETWORK INTERFACE CARDS
DATA MOVER #6
CACHE RAM CARDS
NETWORK INTERFACE CARDS

*Fig. 4*